

Running Head: PARTICLE FILTERS AND RESTLESS BANDITS

Modeling Human Performance in Restless Bandits with Particle Filters

Sheng Kung Michael Yi

Mark Steyvers

Michael Lee

University of California, Irvine

Address for correspondence:

Mark Steyvers

University of California, Irvine

Department of Cognitive Sciences

3151 Social Sciences Plaza

University of California Irvine, CA 92697-5100

E-mail: mark.steyvers@uci.edu

Phone: (949) 824-7642 **Fax:** (949) 824-2307

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 2009		2. REPORT TYPE		3. DATES COVERED 00-00-2009 to 00-00-2009	
4. TITLE AND SUBTITLE Modeling Human Performance in Restless Bandits with Particle Filters				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of California, Irvine, Department of Cognitive Sciences, 3151 Social Sciences Plaza, Irvine, CA, 92697-5100				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES Journal of Problem Solving, 2(2), 33-53					
14. ABSTRACT Bandit problems provide an interesting and widely-used setting for the study of sequential decision-making. In their most basic form, bandit problems require people to choose repeatedly between a small number of alternatives, each of which has an unknown rate of providing reward. We investigate restless bandit problems, where the distributions of reward rates for the alternatives change over time. This dynamic environment encourages the decisionmaker to cycle between states of exploration and exploitation. In one environment we consider the changes occur at discrete, but hidden, time points. In a second environment, changes occur gradually across time. Decision data were collected from people in each environment. Individuals varied substantially in overall performance and the degree to which they switched between alternatives. We modeled human performance in the restless bandit tasks with two particle filter models, one that can approximate the optimal solution to a discrete restless bandit problem, and another simpler particle filter that is more psychologically plausible. It was found that the simple particle filter was able to account for most of the individual differences.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 37	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

Abstract

Bandit problems provide an interesting and widely-used setting for the study of sequential decision-making. In their most basic form, bandit problems require people to choose repeatedly between a small number of alternatives, each of which has an unknown rate of providing reward. We investigate restless bandit problems, where the distributions of reward rates for the alternatives change over time. This dynamic environment encourages the decision-maker to cycle between states of exploration and exploitation. In one environment we consider, the changes occur at discrete, but hidden, time points. In a second environment, changes occur gradually across time. Decision data were collected from people in each environment. Individuals varied substantially in overall performance and the degree to which they switched between alternatives. We modeled human performance in the restless bandit tasks with two particle filter models, one that can approximate the optimal solution to a discrete restless bandit problem, and another simpler particle filter that is more psychologically plausible. It was found that the simple particle filter was able to account for most of the individual differences.

Modeling Human Performance in Restless Bandits with Particle Filters

Many real-world environments involve temporal changes that require decision-makers to adapt their strategies over time. For example, stock market analysts need to track temporal changes in the market carefully, and sport coaches need to track changes in the performance of a team. In some environments, a decision has to be made between different choices, each of which might be associated with uncertain outcomes that can change over time. For example, drivers have to choose between a number of routes each associated with some uncertainty about travel times. In addition, traffic changes can lead to changes in the desirability of routes requiring drivers to adapt their driving strategy continually. In this research, we study how people perform in sequential decision-making situations where each alternative is associated with an uncertain payoff and the underlying environment can change at any time, leading to different payoffs for each alternative.

Bandit problems, as originally described by Robbins (1952), provide a classic task to study sequential decision making. In a standard, stationary bandit environment, people are given a limited number of sequential selections among a fixed set of alternatives, or arms. After each decision, an outcome is generated based on a hidden reward distribution specific to the alternative chosen; the task of the performer is to maximize the total outcomes after all selections have been made. In order to be successful, decision-makers in a bandit environment have to balance their selections between general exploration and exploitation behaviors. Exploration is characterized by the selection of different arms to learn about the hidden outcome distributions for each alternative. Exploitation is characterized by a focus on a single arm, in order to obtain rewards from an option that is believed to be sufficiently good as compared to the other

competing options. An expected behavior in a standard bandit problem may start with a period of exploration, followed by exploitation for the remaining choices.

In a standard bandit problem (also called a ‘game’), the reward rate for each alternative is kept constant over all of the trials. The number of trials in each game may be known, creating a finite horizon problem, or unknown, creating an infinite horizon problem. Optimal solutions can be found for all cases in finite horizon environments using a dynamic programming approach, where optimal decisions are computed for all potential cases starting from the final trial and solving for each trial towards the first (Kaelbling et al., 1996). As the length of a game increases or the number of alternatives increases, the computation necessary to create a complete decision tree increases exponentially. For infinite horizon problems, certain cases may be solved using Gittins indices (Gittins, 1989). A Gittins index gives each alternative a utility that takes into account an alternative’s current estimated value and the information that can be gained from choosing the alternative; the optimal decision is the arm which has the largest index value. However, Gittins indices are only applicable to a limited number of bandit problems, and can be difficult to compute even in those cases (Berry & Fristedt, 1985).

When optimal solutions are available, bandit problems provide an opportunity to examine whether or how people make the best possible decisions. For this reason, many previous empirical studies have been motivated by economic theories, with a focus on deviations from rationality in human decision-making (e.g., Banks, Olson, & Porter, 1997; Meyer & Shi, 1995). More recently, human performance on the bandit problem has been studied within cognitive neuroscience (e.g., Cohen, McClure, & Yu, 2007; Daw, O’Doherty, Dayan, Seymour, & Dolan, 2006) and probabilistic models of human cognition (e.g., Steyvers, Lee, & Wagenmakers, in press).

The environments in empirical studies have ranged from simple, two-choice bandit problems with either one or two non-deterministic arms (Avineri & Prasher, 2006; Banks et al., 1997; Ben-Elia et al. 2008; Meyer & Shi, 1995) to more complicated environments with more than two probabilistic arms (Steyvers et al., in press). While there is evidence of significant variation in how people make decisions, people are able to perform significantly better than chance performance, though few are able to match optimal levels of performance.

Most bandit problem research has focused on stationary bandit problem environments, and there has been relatively little focus on the restless bandit problem, especially in empirical work. In the restless bandit problem, the reward rates for alternatives may change over time, rather than remaining stationary through each trial of a game (Whittle, 1988). The introduction of non-stationary outcome distributions adds a large element of complexity in computing optimal decision processes. But it also provides a strong tie to realistic applications, since most sequential decision-making environments found in real life require consideration of changes in the environment. People making decisions in a restless environment are faced with the additional task of change detection (Brown & Steyvers, 2005; Chinnis & Peterson, 1968, 1970; Massey & Wu, 2005), forcing a continuous switch between exploration and exploitation that is not present in the stationary case.

Few empirical studies have looked at human performance in restless bandit tasks. Estes (1984) looked at human performance in repeated two-armed bandit problem games with one known arm and one fluctuating arm. The known arm provided payoffs at a constant probability, while the fluctuating arm provided payoffs with probabilities in a sine-wave pattern over the course of a block. It was found that subjects made choices in a mild wavelike pattern corresponding to the variation in the alternative reward probabilities. The restless bandit problem

has also been studied via brain imaging by Daw et al. (2006), where brain activity has been found to be correlated to obtained rewards and exploratory decision-making. Finally, there are a few animal learning studies that have measured the ability of animals to adapt to changes in reinforcement schedules. These experiments have shown there are substantial individual differences in the ability to track and respond to changes (Gallistel, 2001).

Here, we use a particle filter approach to finding solutions to the restless bandit problem. Particle filtering is a sequential Monte Carlo method, where a set of particles is updated at each time point to estimate the current state of an environment (see Doucet, de Freitas, & Gordon, 2001). Particles can be thought of as propositions about the environment's current state; as information is gained, particles that describe the observed data well tend to be propagated, while those that do not will tend to be replaced. Over the set of all particles, the propositions form an estimate of the distribution of environment states. These estimates can then be used to make an informed decision on each step of a problem. Particle filters can be used in many situations where Markov Chain Monte Carlo (MCMC) methods become inefficient. For environments where a long history may need to be maintained, the MCMC method will require more computation time with increasing information. In contrast, particle filters, depending on how they are designed, will require less computation time because only a set of hypotheses about the current environmental state needs to be maintained.

Particle filters also hold potential in use as descriptors of human performance (Brown & Steyvers, 2009; Daw & Courville, 2007; Sanborn, Griffiths, & Navarro, 2006). By relaxing or changing model parameters, we obtain behaviors that deviate from the optimal strategy in ways that may be useful in describing human performance on restless bandit problems. An early application of particle filters in cognitive science is provided by Sanborn, Griffiths and Navarro

(2006), who studied sequential effects in category learning. In their modeling, particles correspond to different mental hypotheses about category structures that the human learner might track. By manipulating the number of particles, their category learning model naturally spanned an interesting range of theoretical possibilities. In particular, when restricted to a single particle, their model reduced to Anderson's (1991) classic 'rational model' of category learning, but for a sufficiently large number of particles their model mimicked optimal category learning behavior. In this way, finding the number of particles needed to model people's behavior in sequential category learning tasks provided a natural theoretical mechanism for estimating the complexity of the hypotheses considered by people in learning, and the rationality of their performance.

A second, very recent application of particle filter methods in cognitive science is provided by Brown and Steyvers (2009). These authors applied the particle filter as a descriptor of human performance on an inference and prediction task where the outcome generation distribution changed over time. Individual differences in human performance on the two tasks could be described through shifts in the particle filter's behavior over changes in model parameters. As change detection is a key part of decision-making for the restless bandit, there is potential for application of particle filters to describe individual differences in human behavior for the restless bandit problems as well.

In this paper, we present two different simple restless bandit environments for which particle filter solutions can be employed. We compare these models to the performance of humans in these environments.

Experiment 1

The restless bandit problem used in the first experiment is an extension of sequential stationary infinite-horizon problems. The stationary infinite-horizon bandit problem is one in

which, after each decision trial, there is a set probability γ that the game will end. That is, the distribution of individual game lengths follows a Geometric distribution with parameter γ . While games do have an expected length, there is no way of knowing when a game will end. We can obtain the restless bandit environment by considering the scenario where these infinite-horizon games are played consecutively without breaks, such that the indication of the end of each individual game is removed. Without a clear delineation between the change in reward rates, the decision-maker must have a method for noticing these changes in order to maintain good performance. In the standard bandit problem, the shift between exploration behaviors and exploitation behaviors generally occurs only once in a single game. For the restless bandit, we expect a shift back and forth between exploration and exploitation, as periods of stable reward rates are split by changes that must be detected and accommodated in making decisions.

We observe the behavior of human participants in these restless bandit environments, and compared their performance to two different particle filter methods of solutions. One of these solutions is optimal, while the other sub-optimal but has a more flexible in range of possible behaviors. We compare general behaviors in the sample of human task-performers to that of the models and make observations.

Participants

27 participants drawn from the University of California, Irvine Human Subjects Pool to perform the experimental task for course credit. No demographic information was recorded.

Design

Participants played a series of games (blocks) in restless bandit environments. On each trial t of each game, participants were asked to select one of $N = 4$ alternatives. Each selection D_t

generated an outcome y_t of either a reward or no reward, based on a Bernoulli draw on the selected alternative's reward rate for that trial $\theta_{D_t,t}$ with $y_t \sim \text{Bernoulli}(\theta_{D_t,t})$.

Reward rates θ on each alternative were drawn from a common generating distribution at the start of each game. On each successive trial, the outcome z_t of a Bernoulli draw with parameter γ determined the alternatives' reward rates: with probability $1 - \gamma$, reward rates on all arms on trial t were maintained to be the same as on trial $t-1$, and with probability γ , reward rates on all arms were redrawn from the generating distribution. An example of the outcome generation procedure for a three-arm, 20-trial game can be observed in Figure 2. Arm reward rates are visible in the upper plot, changing at randomly distributed time points. Outcomes generated are visible in the lower plot; the only outcomes that are observed are from the arms that are chosen, denoted by the gray boxes.

Each participant played through a total of $G = 42$ games, each with $K = 50$ trials on $N = 4$ alternatives. In all games, the reward rates θ were drawn from a $\text{Beta}(1,1)$ distribution (i.e., the uniform distribution) and we set the change rate $\gamma = 0.2$. To facilitate the comparison between individual participants and models, a random seed was set such that all participants went through the same sequence of games, with the same potential rewards for each of the alternatives in each game. The first two blocks performed by each participant were excluded from the final analysis as practice blocks; results reflect the actions performed in only the last forty blocks of the experiment.

Apparatus

The task was performed electronically through a program coded in MATLAB. An example of the experiment interface can be seen in Figure 1. Buttons on the right side of the window represented the alternatives to be chosen, and each selection generated an outcome in

the plots to the immediate left of each button. Red bars indicated that no reward was gained, while green bars indicated that a reward was gained. A plot on the far right indicated the number of trials remaining in the current game, while text at the top of the window indicated the current trial, game, and total rewards accumulated. Participant selections and rewards were recorded, as were the hidden reward rates on each trial. Reaction time data were not recorded.

Procedure

Participants were introduced to the task without going into a detailed explanation of the reward-generation process. Participants were simply told that reward rates were randomly generated without describing the precise generating distribution, and that reward rates would change occasionally over the course of each game. Participants were asked to select the alternatives that would maximize their total reward. To maintain focus on the reward maximization objective, a half-second time penalty before the next selection was given if a no reward outcome was generated.

Results

Participant performances over the final forty blocks of the experiment were evaluated with two different measures: the proportion of trials where a reward was obtained; and the proportion of trials where a switch was made. A switch was counted when the alternative chosen in a trial was not the same as that chosen in the previous trial. Figure 3 shows a plot where the performance of each individual participant is marked with a black triangle. It is immediately clear that there is large individual variation in task behavior and performance. A roughly inverted U-shape can be observed, where those with the best performances tend to have a moderate level of switching, while those who switch arms too often or too rarely experience lower performance.

Optimal Particle Filter

An optimal decision-making procedure can be produced by decomposing the problem of solution into two major components. In the first step, based on decisions and outcomes of previous trials, the probability of a change in reward rates is estimated for each trial. In the second step, these change probability estimates can be translated into a distribution of stable periods where there are no changes in reward rate. An optimal decision can be found for each of these periods. By then aggregating over all possible periods and their posterior probabilities of being the true state, the best decision for the next trial of the game can be calculated.

This general method of solution lends itself naturally to computation via Monte Carlo methods. Particle filtering comes as an especially useful way of looking at this restless bandit problem. Each particle contains a single prediction about which trials are associated with a change in reward rates. These predictions also include specify how many trials preceding the current trial are associated with the same reward rates. If we make a prediction of when the next change in reward rate will occur, then we obtain an interval with constant reward rates which turns the problem in a finite bandit problem for which an optimal solution can be obtained. In this particle filter, we can derive optimal solutions for all inferred stable intervals and pick the alternative that corresponds to the mode of the distribution of all optimal solutions. This approximates, in the limit of the number of particles, the optimal decision for a trial.

Particles that better match the observed data have a higher chance of being propagated to future trials, and so over a sufficiently large pool of particles, an accurate estimate of the changepoint distribution can be obtained. With fewer particles, our estimates may be reduced, and the model's performance will worsen. This may, however, be useful in describing human performance and sub-optimality on the task. In the current implementation of the optimal particle filter, we look not only at the optimal case with sufficiently large number of particles and

properly set expected change rate, but also the range of performance over varying numbers of particles and different expected change rates, and compare the optimal model's performance to the range observed in human performance. Details of the optimal particle filter can be found in Appendix A.

Reward Rate Particle Filter

In the optimal particle filter, particles retain estimates of the trials at which changes in the reward rates may have occurred. Each particle can be used to specify a stable interval where no changes in reward rate occur; each interval has a calculable solution and optimal decision. Over multiple particles, the modal alternative that is chosen will be selected as the optimal decision for a trial. While this can provide an optimal solution, this model might not be psychologically plausible. The decision step relies on the ability to compute the optimal decision for finite bandit problems which might unlikely to be available to human performers.

For these reasons, we also considered an alternative approach to solving the problem, which may be more useful in describing human decision-making behavior. In the reward rate particle filter, particles retain an estimate of the reward rate on each alternative. Over all particles, we obtain an estimate of the current reward rate for each alternative in the form of distributions. The decision step is greatly simplified, taking a greedy approach to selection. For each particle, the alternative with the highest reward rate is taken as the best option; the modal alternative chosen over particles is the model decision for the trial. The changes made to the particle filter model reduce the maximum potential performance, but also increase the range of behaviors observable in a fashion that may better account for individual differences observed in human participants. As with the optimal particle filter, we have two parameters we manipulate: number

of particles, and expected change rate. Details of the reward rate particle filter implementation can be found in Appendix B.

Modeling Results

Performance of both the optimal and reward rate particle filters were evaluated in terms of overall reward rate and inter-trial switch probability over the range of model parameters, played over the same forty blocks completed by participants. The range of performance under both particle filter models is marked in Figure 3 by the gray shaded areas. The optimal particle filter's range is plotted in dark gray, and the reward rate particle filter is plotted in light gray.

The optimal particle filter's range of performance clearly does not describe the majority of participants well. The model's reliance on optimal decision behavior in the final decision step brings the base reward rate of the model to that comparable to that of the best human performers, even when the model is limited to maintaining only a single particle at each trial. As the number of particles used is increased, performance also increases, and as the internal estimate of the change rate increases, so do the proportion of trials where a switch is made. For higher internal expectations of change rate, model behaviors are similar to a "Win-Stay, Lose-Shift" strategy: When a reward is obtained on the most recent trial, the same arm is selected with probability p , with a random other arm chosen otherwise; when a reward is not obtained, a random other arm is chosen with probability p , staying on the same arm otherwise. Values of p approaching 1 increase reward gains while reducing the proportion of switch trials.

In comparison to the optimal particle filter, the reward rate particle filter has a much larger range of overall behaviors over the range of parameters. Increases in the number of particles dramatically raises the reward rate, though it asymptotes once the number of particles maintained exceeds 100 to 200. Performance of the reward rate particle filter is comparable to

that of the optimal at the upper limit; the ‘greedy’ strategy of selecting the arm with the highest expected reward rate does not differ significantly in overall reward rate from the strategy employed by the optimal model.

The fact that both the optimal and reward rate particle filters seem to asymptote at such low numbers of particles is interesting. Despite the risk of degeneracy, performance does not suffer from a relatively coarsely-estimated distribution of change points (in the optimal) or current reward rates (in the non-optimal). It requires relatively little effort to obtain a strategy that performs well above random chance; a pure “Win-Stay, Lose-Shift” heuristic strategy performs nearly as well as the particle filter models at peak parameter settings, in terms of overall reward rate. In addition to a ‘greedy’ strategy creating little difference in terms of overall performance, the information that needs to be kept in order to make an informed decision does not have to be particularly large.

Despite the fact that simple strategies can create good performance, the fact remains that our empirical results show human decision-makers follow a wide range of behaviors, and achieving varying degrees of success. Heuristic strategies and optimal decision-making models are too narrow to account for these individual differences well. The sub-optimal particle filter model, however, perhaps has the flexibility to describe individual performances with intuitively interpretable parameters, while also maintaining the ability to perform at near optimal levels with proper parameter choices.

Experiment 2

In Experiment 1, we found that the performance of the reward rate particle filter was, over the range of parameter values, more adept at describing a larger variety of potential behaviors, including those of human subjects, than that of an optimally designed particle filter. In

addition, there was relatively little loss in maximum expected rewards when using the reward rate particle filter as compared to the optimal. The reward rate particle filter also carries the advantage that the decision step is much simpler and the information contained in individual particles does not require a memory of previous trials. These properties give the reward rate particle filter the potential for applicability in a wider array of environments than the one for which it was originally designed.

For Experiment 2, we look at a different restless bandit environment where changes occur continuously over time, rather than at randomly spaced intervals. On each trial, the reward probability on each arm is re-sampled from a distribution centered about the reward probability from the previous trial. Thus, while there is expected to be little change in reward rate from trial to trial, over multiple trials, the reward rates may change drastically. In this environment, the best strategy to employ at each time point is to choose the alternative with the largest estimated reward rate.

Once again, we observe the behavior of human decision-makers in this environment and compare them to particle filter models. Here, we employ two types of reward rate particle filters: one that has a propagation mechanism that matches the environment and a second identical to the one used in Experiment 1 that does not match the environment. Our goal is to observe how much loss the use of the inappropriate reward rate particle filter incurs, and how it compares to the optimal reward rate particle filter.

Participants

36 participants were drawn from the University of California, Irvine Human Subjects Pool to perform the experimental task, for course credit. No demographic information was recorded.

Design

All aspects of the experimental design were maintained from Experiment 1 except for the reward rate generation mechanism, such that changes in reward rates occurred continuously over trials, rather than at discrete time points. On the first trial of each game, reward rates were drawn from a common generating distribution $Beta(1,1)$, as in Experiment 1. On each subsequent trial, a new reward rate was drawn on each arm dependent on the previous reward rate, $\theta_{i,t} \sim Beta(1+c(\theta_{i,t-1}), 1+c(1-\theta_{i,t-1}))$, where c is a parameter that controls the variability in reward rate between trials. The mean of the distribution, $\frac{1+c(\theta_{i,t-1})}{2+c}$, is slightly biased toward the neutral value of 0.5. Increasing parameter c decreases the variance in the generating distribution and shifts the mean toward that of the previous trial's reward rate. An example of a three-arm, 20-trial game can be observed in Figure 4. As with Figure 2, arm rates are observable in the upper panel, while outcomes can be seen in the lower panel.

As with Experiment 1, each participant played through the same sequence of $G = 42$ games, each with $K = 50$ trials and $N = 4$ alternatives with $c = 10$. The first two blocks performed by each participant were excluded from the final analysis as practice blocks; results reflect the actions performed in only the last forty blocks of the experiment.

Apparatus

Aside from the reward rate generation method, no changes were made to the experimental program.

Procedure

The procedure was nearly identical to that of Experiment 1. Only a slight change to experimental instructions was made to reflect the changes in reward rate distribution; changes were stated to occur gradually throughout each game, rather than occasionally at random times.

Results

Participant performance was again evaluated on the final forty blocks played, in terms of the proportion of trials where a reward was obtained, and the proportion of trials where a switch was made. Figure 5 shows participant performance on these measures, marked with black triangles. As with Experiment 1, there are large individual differences in participant behavior. The inverse U-pattern observed in the previous experiment is again visible, although it is not quite as clear. A large group of participants with switch proportions between approximately 0.3 and 0.45 have overall behaviors that are in line with the “Win-Stay, Lose-Shift” strategy.

Continual-change Reward Rate Particle Filter

With the change in the reward-generation environment, changes in decision-making strategy must also be made to maintain good performance. A particle filter in which particles carry information regarding reward rates should have each particle, when propagated, adjust its reward rate estimates in the same way as the environment. Sampled particles’ estimated reward rates will first be re-sampled from distributions estimating the level of expected drift in trial-wise reward rate before being potentially propagated to the next trial. The decision step of the adjusted particle filter remains the same as before, taking the alternative with the highest estimated reward rate. This particle filter has two parameters that can be varied. The first, the number of particles, serves the same purpose as in the models for Experiment 1. The second parameter, representing the expected variability in reward rates between trials, serves a similar function to the expected change rate in the Experiment 1 models. Details of the second reward rate particle filter can be found in Appendix C.

Modeling Results

Performance measures were obtained for the continual-change and discrete-change reward rate particle filter models across model parameters in terms of reward rate and switch rate. Figures 5 and 6 show the range of performance under both particle filter models, marked by the gray shaded areas. Figure 5 includes both the discrete-change (dark gray) and continual-change (light gray) particle filters, while Figure 6 includes the discrete-change particle filter alone.

The continual-change reward rate particle filter model shows a wide range of overall behaviors across its parameters. Performance changes in an expected pattern: increases in the number of particles result in an increase in the overall reward rate; increases in the expected variability of the environment reward rates results in an increase in switching between alternatives. The discrete-change particle filter shows the same pattern as in Experiment 1, though flattened both due to the difficulty of the continuous-change environment as well as the mismatch between propagation method and environment generation. Comparing the performance of the discrete-change particle filter to that of the continual-change particle filter, we find that at the highest particle counts, the continual-change particle filter does markedly better than the discrete-change model. As a result, there are participants with lower switch rates and high reward rates who are not well-described by the discrete-change particle filter, but fall under the range covered by the continual-change model.

We also compare the continual-change model to the discrete-change model in the environment of Experiment 1. Figure 7 plots the discrete-change (dark gray) and continual-change (light gray) particle filter together in the discrete change-point environment across the range of model parameters. While in Experiment 2, the discrete-change particle filter lost performance due to environment change, there is remarkably little difference between the two models' overall performance range in the Experiment 1 environment.

Discussion

Bandit problems have been utilized extensively in sequential decision-making research, but relatively little empirical research has been done with restless environments, where reward rates may change over time. Here, we have observed the performance of people in two restless bandit environments with different reward rate change dynamics. In the first, reward rates changed at discrete, but random time points, for all arms simultaneously; in the second, reward rates changed continuously such that there were small short-term changes and large long-term changes. We found that people were able to perform the task in both environments but we also found substantial individual differences in behavior. Generally, participants who performed best were those who switched options at an appropriate rate, while those who switched too much or too little performed comparatively worse. This range of behaviors was well described by particle filter models. Each the particle filter that we considered had easily-interpretable parameters. One parameter is the number of particles that modulates the overall performance. The number of particles can be likened to the amount of cognitive resources applied to solving a problem. As the number of particles increases, performance increases with diminishing returns for large numbers of particles. The other parameter related to the perceived variability in the environment and modulated the amount of switching between alternatives. We found that varying these parameters explained most of the observed individual differences.

A natural application of the particle filter approach would be to assess the best-fitting parameters for individual participants. This would allow a natural explanation of their behavior in terms of interpretable parameters. However, the probabilistic nature of the particle filter model creates a lot of variability in overall performance, making it difficult to obtain a precise estimate of the average behavior for a specific parameter pair. Thus, it is hard to give a precise

characterization of a person's performance in terms of a best-fitting particle filter model. While there is a clear relationship between parameter values and behavior, and parameters do have easily interpretable meanings, there is quite a bit of variability in what can be expected from the model's performance.

Overall, particle filters provide a promising modeling framework not only to approximate the optimal solution in restless bandit environments but also to model individual differences.

Acknowledgments

This work was funded by award FA9550-07-1-0082 from the Air Force Office of Scientific Research.

References

- Anderson, J. R. (1991). The adaptive nature of human categorization. *Psychological Review*, 98(3):409-429.
- Avineri, E. & Prashker, J. N. (2006) The impact of travel time information on travelers' learning under uncertainty. *Transportation*, 33(4), 393-408.
- Banks, J., Olson, M., & Porter, D. (1997) An experimental analysis of the bandit problem. *Economic Theory*, 10, 55–77.
- Ben-Elia, E., Erev, I., Shiftan, Y. (2007) The combined effect of information and experience on drivers' route-choice behavior. *Transportation*, 35(2), 165-177.
- Berry, D. A., & Fristedt, B. (1985) *Bandit Problems: Sequential Allocation of Experiments*. New York, Chapman and Hall.
- Biele, G., Erev, I., Ert, E. (in press). Learning, risk attitude and hot stoves in restless bandit problems. *Journal of Mathematical Psychology*.
- Brown, S., Steyvers, M. (2009). Detecting and predicting changes. *Cognitive Psychology*, 58, 49-67.
- Cohen, J. D., McClure, S. M., & Yu, A. J. (2007). Should I stay or should I go? exploration versus exploitation. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362, 933–942.
- Chinnis, J., & Peterson, C. (1968) Inference About a Nonstationary Process. *Journal of Experimental Psychology*, 77, 620-625.
- Chinnis, J., & Peterson, C. (1970) Nonstationary Processes and Conservative Inference. *Journal of Experimental Psychology*, 84, 248-251.

Daw, N. D., & Courville, A. C. (2007). The rat as particle filter. *Advances in Neural Information Processing*, 20, 369–376.

Daw, N. D., O’Doherty, J. P., Dayan, P., Seymour, B., & Dolan, R. J. (2006) Cortical substrates for exploratory decisions in humans. *Nature*, 441, 876–879.

Denrell, J., & March, J. G. (2001) Adaptation as Information Restriction: The Hot Stove Effect. *Organization Science*, 12(5), 523-538.

Doucet, A., de Freitas, N., & Gordon, N. (2001). Sequential Monte Carlo methods in practice. New York: Springer.

Estes, W.K. (1984) Global and Local Control of Choice Behavior by Cyclically Varying Outcome Probabilities. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 10(2), 258-270.

Gallistel, C. R., Mark, T., King, A., & Latham, P. E. (2001). The rat approximates an ideal detector of changes in rates of reward: Implications for the law of effect. *Journal of Experimental Psychology: Animal Behavior Processes*, 27, 354–372.

Gittins, J. C. (1989). *Multi-armed bandit allocation indices*. New York: Wiley.

Kaelbling, L. P., Littman, M. L., Moore, A. W. (1996) Reinforcement Learning: A Survey. *Journal of Artificial Intelligence Research*, 4, 237-285.

Massey, C., Wu, G. (2005) Detecting Regime Shifts: The Causes of Under- and Overreaction. *Management Science*, 51(6), 932-947.

Meyer, R. J., Shi, Y. (1995) Sequential Choice Under Ambiguity: Intuitive Solutions to the Armed Bandit Problem. *Management Science*, 41(5), 817-834.

Robbins, H. (1952). Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society*, 55, 527–535.

Sanborn, A. N., Griffiths, T. L., & Navarro, D. J. (2006). A more rational model of categorization. *Proceedings of the 28th Annual Conference of the Cognitive Science Society*.

Steyvers, M., Lee, M.D., & Wagenmakers, E.J. (in press) A Bayesian analysis of human decision-making on bandit problems. *Journal of Mathematical Psychology*.

Whittle, P. (1988). Restless Bandits: Activity Allocation in a Changing World. *Journal of Applied Probability*, 25A, pp. 287–298.

Appendix A

Details of Optimal Particle Filter for Experiment 1

The optimal particle filter solution to the given restless bandit environment follows two major phases. In the first phase, we propagate particles from the previous trial to the current trial, after taking into account the most recently observed outcome. The current implementation of the particle filter uses a direct simulation method, using the following steps:

For $t = 1$, initialize particles $k = 1, \dots, P$, $\mathbf{z}_1^k \sim \text{Bernoulli}(\gamma)$ for each arm, where γ is the expected probability of a change in reward rates.

For $t = 2, \dots, K$,

Initialize counter $p = 0$.

While $p < P$,

Take sample $k \sim U[1, \dots, P]$.

Generate proposal $\hat{\mathbf{z}} = \{\mathbf{z}_{t-1}^k, z_t \sim \text{Bernoulli}(\gamma)\}$ on each arm.

Sample $u \sim U[0, 1]$.

If $P(y_t | \hat{\mathbf{z}}) > u$, $p = p + 1$; $\mathbf{z}_t^p = \hat{\mathbf{z}}$.

In the second phase, each particle is used to describe an interval where no changes in reward rates occur that includes the current trial. The trial of the most recent change point forms the start of the interval, while a draw from a *Geometric*(γ) specifies how many trials remain before the end of the period. These intervals describe finite-horizon bandit problems, whose solutions can be obtained by dynamic programming (Kaelbling et al. 1996). Each interval has an optimal selection for the current trial; the mode selection over all particles is the optimal choice

for the model. Two parameters can be manipulated, the number of particles P that are maintained on each trial, and the expected change rate γ .

Appendix B

Details of the Discrete-Change Reward Rate Particle Filter

As with the optimal particle filter, each trial's decision is based on a two-step solution of propagating particles, then selecting an alternative based on the mode response over particles. The first step is very similar to that of the optimal particle filter, propagating particles through a direct simulation method, except that particles carry reward rate information at the most recent trial, rather than change point information over the game played so far:

For $t = 1$, initialize particles $k = 1, \dots, P$, $\theta_1^k \sim \text{Beta}(1,1)$ for each arm

For $t = 2, \dots, K$,

Initialize counter $p = 0$.

While $p < P$,

Take sample $k \sim U[1, \dots, P]$.

Generate proposal $\hat{\theta}$ on each arm:

If $\text{Bernoulli}(\gamma) = 1$, $\hat{\theta} \sim \text{Beta}(1,1)$ for each arm,

Otherwise $\hat{\theta} = \theta_{t-1}^k$.

Sample $u \sim U[0,1]$.

If $P(y_t | \hat{\theta}) > u$, $p = p + 1$; $\theta_t^p = \hat{\theta}$.

The decision step is considerably simpler than in the optimal model. Each particle gives a predicted reward rate on each arm; the best decision implied by each particle is the alternative with the largest reward rate. The mode alternative over all particles is the choice made by the model for the next presented trial. As with the optimal particle filter, we can manipulate two

parameters, the number of particles P that are maintained on each trial, and the expected change rate γ .

Appendix C

Details of Continual-Change Rate Particle Filter

The continual-change rate particle filter operates similarly to the discrete change particle filter, with an identical decision method. What differs is the particle propagation method - after a particle is sampled, the proposal particle reward rates are sampled from a distribution centered around the reward rates of the sampled particle:

For $t = 1$, initialize particles $k = 1, \dots, P$, $\theta_1^k \sim \text{Beta}(1, 1)$

For $t = 2, \dots, K$,

Initialize counter $p = 0$.

While $p < P$,

Take sample $k \sim U[1, \dots, P]$.

Generate proposal $\hat{\theta}$ on each arm:

$$\hat{\theta}_i \sim \text{Beta}(1 + c\hat{\theta}_i, 1 + c(1 - \hat{\theta}_i)) \text{ for } i = 1, \dots, N.$$

Sample $u \sim U[0, 1]$.

If $P(y_t | \hat{\theta}) > u$, $p = p + 1$; $\theta_t^p = \hat{\theta}$.

As with the discrete-change reward rate particle filter, the alternative chosen at the decision step is the mode decision over all particles, where each particle implies the best decision to be on the arm with the largest reward rate. There are two parameters that we can manipulate, the number of particles P that are maintained on each trial, and the estimate of variability in reward rates between trials c .

Figure Captions

Figure 1. Sample task interface for Experiment 1 and 2.

Figure 2. Sample of outcome generation for Experiment 1. Hidden reward rates change at discrete time points and are redrawn from the generating distribution at each change. Outcomes from each arm are generated based on reward rates; only outcomes from selected arms are made visible.

Figure 3. Subject performances in Experiment 1 (black triangles), against the range of the optimal particle filter (dark gray) and reward rate particle filter (light gray).

Figure 4. Sample of outcome generation for Experiment 2. Hidden reward rates change after each trial, based on the reward rate on the previous trial. Outcomes from each arm are generated based on reward rates; only outcomes from selected arms are made visible.

Figure 5. Subject performances in Experiment 2 (black triangles) against the range of the continual-change reward rate particle filter (dark gray) and discrete-change reward rate particle filter (light gray).

Figure 6. Subject performances in Experiment 2 (black triangles) against the range of the discrete-change reward rate particle filter (light gray).


Figure 7. Subject performances in Experiment 1 (black triangles) against the range of the discrete-change reward rate particle filter (light gray) and continual-change reward rate particle filter (dark gray).

Trial 25 of 50

Game 4 of 40

Points this game: 14

Total Points Earned: 101



- 1
- 2
- 3
- 4



